

FREQUENCY EXTRAPOLATION BY NONCONVEX COMPRESSIVE SENSING

*Rick Chartrand**

Los Alamos National Laboratory
Theoretical Division, MS B284
Los Alamos, NM 87545

Emil Y. Sidky and Xiaochuan Pan†

University of Chicago
Department of Radiology
5841 S. Maryland Ave., Chicago, IL 60637

ABSTRACT

Tomographic imaging modalities sample subjects with a discrete, finite set of measurements, while the underlying object function is continuous. Because of this, inversion of the imaging model, even under ideal conditions, necessarily entails approximation. The error incurred by this approximation can be important when there is rapid variation in the object function or when the objects of interest are small. In this work, we investigate this issue with the Fourier transform (FT), which can be taken as the imaging model for magnetic resonance imaging (MRI) or some forms of wave imaging. Compressive sensing has been successful for inverting this data model when only a sparse set of samples are available. We apply the compressive sensing principle to a somewhat related problem of frequency extrapolation, where the object function is represented by a super-resolution grid with many more pixels than FT measurements. The image on the super-resolution grid is obtained through nonconvex minimization. The method fully utilizes the available FT samples, while controlling aliasing and ringing. The algorithm is demonstrated with continuous FT samples of the Shepp-Logan phantom with additional small, high-contrast objects.

Index Terms— Frequency extrapolation, aliasing, Fourier transform imaging, MRI, compressive sensing

1. INTRODUCTION

Tomographic imaging strives to reconstruct digital images of continuous objects from a finite set of discrete measurements. In MRI, for example, the free induction decay samples are interpreted as spatial frequency samples of an underlying object function, related to the proton spin density. An MRI pulse sequence can be designed to obtain a Cartesian grid of spatial frequencies. This grid can efficiently be processed into a discrete image array by the fast Fourier transform (FFT),

which performs an inverse discrete Fourier transform (FT) on the data. Even neglecting physical factors and noise that spoil the imaging model, strong imaging artifacts can appear due to aliasing and ringing. These artifacts are particularly disturbing when viewing low-contrast features in the neighborhood of a large discontinuity or when imaging small objects. Traditional methods of dealing with such artifacts generally involve filtering the frequency values so that the high frequency components are brought gently to zero. Such methods are effective at removing ringing artifacts, but they also throw out the information contained in the frequency components that were attenuated. This loss of information can be particularly disturbing for small objects. A preferable approach is to extrapolate the available measurements to higher frequencies.

Although frequency extrapolation is a long-standing problem in sampling theory [1], we have been motivated to examine this issue again due to the recent interest in compressive sensing (CS). CS methods seek to recover an image from a sparse set of measurements. The issue of frequency extrapolation becomes relevant to the CS community when such algorithms are applied to actual scanner data. Most work in CS begins with a discrete-to-discrete data model, which is further removed from reality than the continuous-to-discrete model for a scanner. (A notable exception is the Xampling framework, which does connect back to a continuous signal model [2]). Take the discrete 2D FT as the imaging model, for example. If a Cartesian grid of frequency samples with dimension N^2 is measured, a discrete image is obtained through the inverse discrete FT. In this case, image recovery from sparse samples is well defined; reconstruction of the N^2 discrete image from a number of samples less than N^2 can be interpreted as sparse data reconstruction.

For continuous-to-discrete imaging models, it is no longer clear what constitutes a sparsely-sampled or fully-sampled imaging system. Arguments based on the Nyquist frequency are generally not applicable for tomography, because object functions are not band-limited, so in principle spatial frequency information is needed over the whole 2D plane. Thought of in another way, CS has been primarily concerned with interpolation of spatial frequencies from a finite set of measurements. But what is needed for CS inversion of the

*This work was supported in part by the U.S. Department of Energy through the LANL/LDRD Program.

†This work was supported in part by NIH grants R01CA120540, R01EB000225 and R01CA112163. The contents of this paper are solely the responsibility of the authors and do not necessarily represent the official NIH view.

continuous FT is a combination of frequency interpolation *and extrapolation*. In this article, we address the latter. Interestingly, the CS principle may provide useful algorithms, which can effectively address frequency extrapolation for certain classes of object functions. And since nonconvex CS has been shown to give better results for underdetermined problems [3, 4, 5], we investigate whether nonconvex methods will improve frequency extrapolation.

In Section 2 we specify the frequency extrapolation problem; in Section 3 we summarize the nonconvex CS algorithm; and in Section 4 we show results of the algorithm applied to FT data from the continuous Shepp-Logan phantom with small test objects embedded.

2. DIFFICULTIES IN INVERTING THE CONTINUOUS-TO-DISCRETE IMAGING MODEL

We demonstrate the frequency extrapolation problem with the inversion of a 512×512 grid of FT samples computed from the continuous FT of the Shepp-Logan phantom with additional test objects. Shown in Figure 1 is the phantom in a low-contrast window over the whole object and in a high-contrast window zooming in on the small test objects. The result of the inverse discrete FT of the data is shown in Figure 2. Clear aliasing artifacts are visible, and the image of the small objects is quite distorted. That the small objects are distorted is no surprise as they occupy only a few pixels, and the pixel representation is clearly not going to be accurate at this length scale. On the other hand, small objects such as these can be quite significant in many imaging applications or small details on larger objects, such as spiculations extending from a malignant tumor, can be important for clinical applications.




Fig. 1. Left: High-resolution, 4096×4096 discretization of a continuous Shepp-Logan phantom, with very small test objects embedded towards the upper right. Right: a 256×256 zoom on the test objects, with an adjusted grayscale window to show contrast.

A common strategy to image such small objects is to employ a higher resolution grid in the image space or, equivalently, zero-pad the frequency samples. Figure 3 shows the




Fig. 2. Inverse DFT of the 512×512 samples of the Fourier transform of the continuous phantom, with a zoom-in (now 32×32) on the small objects. Aliasing is prevalent.




Fig. 3. Inverse DFT of the 4096×4096 array obtained by zero padding the 512×512 FT samples. Ringing is substantial.

result of frequency zero-padding by a factor of 8 in both dimensions. While the small object representation is smoother, there are still conspicuous ringing artifacts. It is well known that these artifacts are caused by the discontinuous jump to zero at the boundary of the original frequency samples. A common procedure is to roll off the high frequency samples in order to eliminate or reduce the size of these discontinuities. There are many ways to do this, and Figure 4 shows how Gaussian filtering removes ringing artifacts from the zero-padding result. Though effective, such filtering necessarily throws out information and important shape detail is lost.

The issue at hand is really one of frequency extrapolation. The inverse DFT assumes that the frequency samples repeat, while zero-padding assumes the unknown frequencies are zero. Both strategies are highly unrealistic for this phantom, as seen in the expanded set of 4096×4096 FT samples in Figure 7. More advanced approaches involve formulating frequency extrapolation as an optimization problem. For example, the Papoulis-Gerchberg algorithm [1] enforces compact support of the image. The same reference also describes an algorithm that is based on information theory and calls for maximizing entropy. The information theory argument is interesting in that it is a reaction to zero padding—a condition that is strong and wrong. Maximum entropy methods aim




Fig. 4. Inverse DFT of the 4096×4096 array obtained by Gaussian filtering the zero-padded FT samples. Aliasing is removed but the small objects are heavily blurred.

to extrapolate to high frequency values that are maximally non-committal [6]. Strong prior information, however, is only detrimental if it is wrong.

The CS principle provides a strong prior based on some form of sparsity in the image. We have investigated sparsity of the image gradient in applications with real X-ray computed tomography data [7, 8], and have found it effective at reducing the necessary number of views as measured quantitatively by many image quality metrics [9]. This same CS principle should be effective when applied to the present frequency extrapolation, as the phantom has a sparse gradient. Unlike other CS applications involving Fourier sampling [3, 4], we do not expect exact recovery even under these ideal conditions for two reasons: (1) although we have spatial sparsity, it still requires an infinite number of infinitesimally small pixels to represent the edges, and (2) the frequency extrapolation problem is more unstable than interpolation. The goal instead is to be able to exploit all the available frequency information to see details down to the minimum length scale of the imaging device. The image quality test here is visual. The small test objects are simple geometric shapes, and the goal is to be able to visually resolve these shapes. Future work will formulate this task as a quantitative image quality metric.

We do not claim that the CS principle provides a frequency extrapolation algorithm that is superior to other algorithms for all objects. Instead, it may be useful for certain classes of objects. Furthermore, this investigation may provide insight into more traditional CS (if CS can be considered traditional) configurations, where both frequency interpolation and extrapolation are necessary.

3. NONCONVEX CS FOR FREQUENCY EXTRAPOLATION

We attempt to reconstruct a 4096×4096 image from the 512×512 grid of frequency samples of our continuous test object. The essence of our approach is to find the image having the sparsest gradient that is consistent with the data. This would

correspond to the following optimization problem:

$$x^* = \underset{x}{\operatorname{argmin}} \|\nabla x\|_0, \text{ subject to } Ax = b. \quad (1)$$

Here, ∇ is a discretization of the gradient operator (for which we use simple forward differences, and periodic boundary conditions); $\|\cdot\|_0$ counts the number of nonzero components; and A is the rectangular DFT operator that transforms an image on the super-resolution, 4096×4096 grid to the FT samples b on the original, 512×512 frequency grid. Note that the linear system $Ax = b$ is severely underdetermined, as each x has 64 times as many elements as b .

The usual CS approach [10] is to replace the intractable problem (1) with a more readily solved optimization problem, using a sparsity-inducing objective function. This is most commonly the ℓ^1 norm; however, it has been shown [3, 4, 5] that better results can be obtained using a *nonconvex* objective function instead. We use an approach similar to that of [11], which uses a penalty function φ_p that is a shifted and scaled p th power function, spliced smoothly with a quadratic function near the origin (see Figure 5). Using $p < 1$ allows the image gradient to be less sparse for the same amount of data, and can preserve shapes better [12].




Fig. 5. Penalty function applied to each pixel of the image gradient, for three values of p . Using $p < 1$ results in sublinear growth, while using $p < 0$ results in $\varphi_p(t)$ being bounded above. This gives a better approximation to the counting norm, which is 0 at $t = 0$ and 1 elsewhere.

Because there is not a perfect match between the DFT of the discrete image and the FT samples of the continuous phantom, we relax the data equality constraint in (1). We use a data fidelity term instead, which penalizes but does not prohibit discrepancy from the data. This gives us the following, unconstrained optimization problem:

$$x^* = \underset{x}{\operatorname{argmin}} \sum_{i=1}^N \varphi_p(|\nabla x|_i) + \lambda \|Ax - b\|_2^2, \quad (2)$$

where $N = 4096^2$ is the number of pixels, and λ is a parameter determining the balance between the competing effects of regularization and data fidelity. For this work λ is chosen by visual inspection. Future work will explore dependence on λ of quantitative image-quality metrics.

To solve (2), we use a modification of the algorithm of [11], which uses a Douglas-Rachford splitting, augmented Lagrangian approach (cf. [13, 14, 15] for the convex case). This decomposes the optimization task into two simpler sub-problems, one requiring a simple generalization of shrinkage (or soft thresholding), the other requiring little more than two FFTs. This makes the algorithm very efficient, which makes it feasible for a reconstruction problem of this size. The modification simply removes one of the two augmented Lagrangian steps, which was designed to enforce the data constraint exactly. Further details are in [11].

4. RESULTS

We demonstrate (2) on the frequency extrapolation problem, for values of $p = 1$, $1/4$, and $-1/2$. The reconstructions for these cases are shown in Figure 6. We note that the visually optimal choice of λ in (2) can depend on the structures being imaged. We choose the largest value (or weakest regularization) that removes nearly all of the aliasing near the small objects. For the phantom as a whole, a tighter data constraint would be appropriate.

When $p = 1$, the small objects are blurred, though less than in the case of filtered, zero-padded data. The blurring is greatly decreased in the $p = 1/4$ case, while with $p = -1/2$ the edges are sharp. These two cases have blockiness in the arms of the crescent, which in the $p = 1$ case is a smooth decrease in intensity. The rectangle is more rounded off in the $p = -1/2$ case, which also struggled to remove all the aliasing without over-regularizing the phantom. The $p = 1/4$ case seems best overall, though for applications where sharp edges are very important, the $p = -1/2$ case would be best.

In order to see how the algorithm performs on its designed function of frequency extrapolation, we compute the extrapolated values for each p by computing the DFT of the reconstruction, and compare with the true values from the analytically computed FT of the continuous phantom. The results in Figure 7 shows that the smaller the value of p , the better the extrapolation, though in all cases the result is substantially more accurate than zero padding.

5. SUMMARY AND CONCLUSION

We have applied the recent principle of compressive sensing to the mature signal processing problem of frequency extrapolation. The corresponding nonconvex CS algorithm appears to be effective for revealing small details in the subject for a class of objects that are approximately piecewise constant. The gain in image quality over basic filtering methods comes from the utilization of all measured frequency components. The present frequency extrapolation method can be applied directly to current FT imaging modalities such as MRI, and it may allow for better image quality with currently used image acquisition schemes.




Fig. 6. Our 4096×4096 reconstructions from 512×512 FT samples of the continuous phantom, using (2) with (from top to bottom) $p = 1$, $1/4$, and $-1/2$. The smaller the p , the sharper the edges, though $p = -1/2$ rounds off the rectangle more than $p = 1/4$, and also over-regularizes the phantom.

Looking down the road, the frequency extrapolation problem will become a necessary part of CS as the focus of CS research shifts to applications. Of course there may be many alternate strategies: [16] promotes a method where the support of the non-zero components of the signal is reconstructed first and the signal expansion is performed only within this support; and [17] describes an algorithm for low-dose X-ray computed tomography, first performing a frequency extrapolation of the CT projection data followed by interpolation of the view angles through a CS approach. Our initial experience in attempting to improve resolution properties is that there will need to be many alternative approaches, because the details of the frequency extrapolation will depend strongly on the imaging model of a particular tomographic system.




Fig. 7. Left to right: actual 4096×4096 FT samples of the continuous phantom, scaled by $|\nu|^{3/2}$ to make high-frequency components visible, with a box showing the portion containing the original 512×512 data; and the DFT of the 4096×4096 reconstructions from the 512×512 FT samples using (2) using $p = 1, 1/4$, and $-1/2$, all scaled the same way and on the same (false) color scale. The smaller the value of p , the better the frequency extrapolation effect.

6. REFERENCES

- [1] A. Papoulis, "A new algorithm in spectral analysis and band-limited extrapolation," *IEEE Trans. Circuits Sys.*, vol. 22, no. 9, pp. 735–742, 1975.
- [2] M. Mishali and Y. C. Eldar, "From theory to practice: Sub-Nyquist sampling of sparse wideband analog signals," *IEEE J. Sel. Topics Signal Process.*, vol. 4, pp. 375–391, 2010.
- [3] Rick Chartrand, "Exact reconstructions of sparse signals via nonconvex minimization," *IEEE Signal Process. Lett.*, vol. 14, pp. 707–710, 2007.
- [4] Rick Chartrand, "Nonconvex compressive sensing and reconstruction of gradient-sparse images: random vs. tomographic Fourier sampling," in *IEEE International Conference on Image Processing (ICIP)*, 2008.
- [5] Rick Chartrand and Valentina Staneva, "Restricted isometry properties and nonconvex compressive sensing," *Inverse Problems*, vol. 24, no. 035020, pp. 1–14, 2008.
- [6] H. H. Barrett and K. J. Myers, *Foundations of Image Science*, John Wiley & Sons, Hoboken, NJ, 2004.
- [7] E. Y. Sidky, C.-M. Kao, and X. Pan, "Accurate image reconstruction from few-views and limited-angle data in divergent-beam CT," *J. X-ray Sci. Tech.*, vol. 14, pp. 119–139, 2006.
- [8] E. Y. Sidky and X. Pan, "Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization," *Phys. Med. Biol.*, vol. 53, pp. 4777–4807, 2008.
- [9] J. Bian, J. H. Siewerdsen, X. Han, E. Y. Sidky, J. L. Prince, C. A. Pelizzari, and X. Pan, "Evaluation of sparse-view reconstruction from flat-panel-detector cone-beam CT," *Phys. Med. Biol.*, vol. 55, pp. 6575–6600, 2010.
- [10] Emmanuel J. Candès, Justin Romberg, and Terence Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, 2006.
- [11] Rick Chartrand, "Fast algorithms for nonconvex compressive sensing: MRI reconstruction from very few data," in *IEEE International Symposium on Biomedical Imaging*, 2009.
- [12] Rick Chartrand, "Nonconvex regularization for shape preservation," in *IEEE International Conference on Image Processing (ICIP)*, 2007.
- [13] Junfeng Yang, Wotao Yin, Yin Zhang, and Yilun Wang, "A fast algorithm for edge-preserving variational multi-channel image restoration," *SIAM J. Imaging Sci.*, vol. 2, pp. 569–592, 2009.
- [14] Tom Goldstein and Stanley Osher, "The split Bregman method for L1 regularized problems," *SIAM J. Imaging Sci.*, vol. 2, pp. 323–343, 2009.
- [15] Simon Setzer, "Split Bregman algorithm, Douglas-Rachford splitting and frame shrinkage," in *Scale Space and Variational Methods in Computer Vision*, vol. 5567 of *Lecture Notes in Computer Science*, pp. 464–476. Springer Berlin / Heidelberg, 2009.
- [16] M. Mishali and Y. C. Eldar, "Reduce and boost: Recovering arbitrary sets of jointly sparse vectors," *IEEE Trans. Signal Process.*, vol. 56, pp. 4692–4702, 2008.
- [17] E. Y. Sidky, Y. Duchin, C. Ullberg, and X. Pan, "A constrained, total-variation minimization algorithm for low-intensity X-ray CT," Submitted.